

## How Rings assign address space

Step1: align incoming address to self (to some power of 2)

Step2: assign the result to self address

Step3:  $\text{next\_addr} = \text{self\_addr} + \text{self\_addr\_space}$ ; // number of register used locally


Step4: send down next\_addr

### Example:

Dma needs 16 addr

Uart needs 4

Timer needs 256



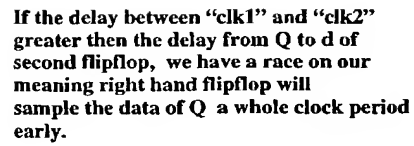

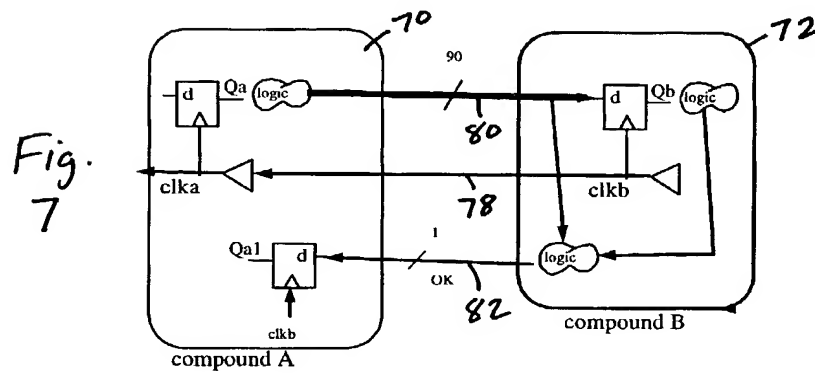
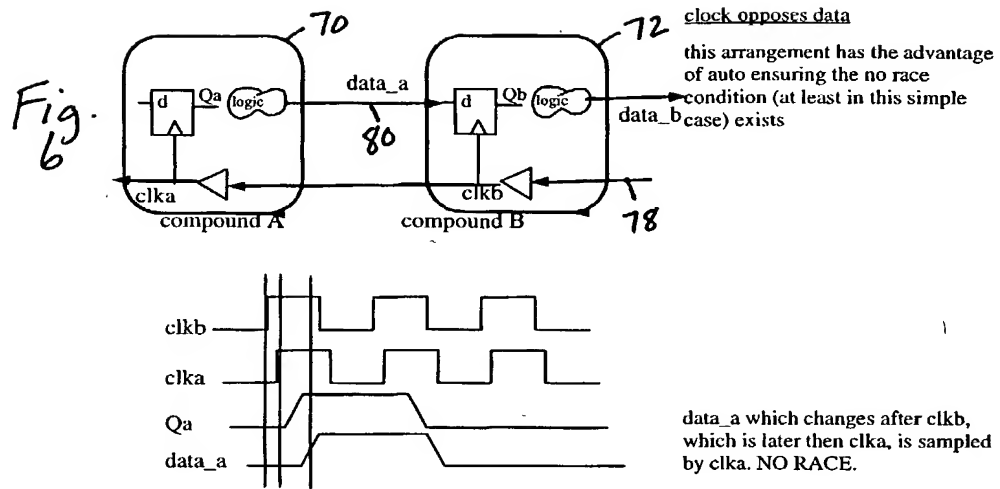


Fig.  
5



the problem is possible race. However, we control the logic on each flipflop leaving the compound, because it is always the same standard ring-interface module. we can ensure, that the delay will be at least enough. And more importantly easily checked after layout.





Fig.  
8

Fig. 9

data\_a leaving the bridge goes to member "b" and there should be sampled by rising of clkb. clkb lags a lot behind clka of the bridge. As clearly seen from the waveforms, race is eminent. Here we should add latches for all the data lines (~90). Adding latch works however if the delay between clka and clkb is less than 75% of cycle time. otherwise the uncertainty kills the usable time. It sets hard limit on the number of ring members. Also keep in mind that latches needed on each OK signal between members of the ring

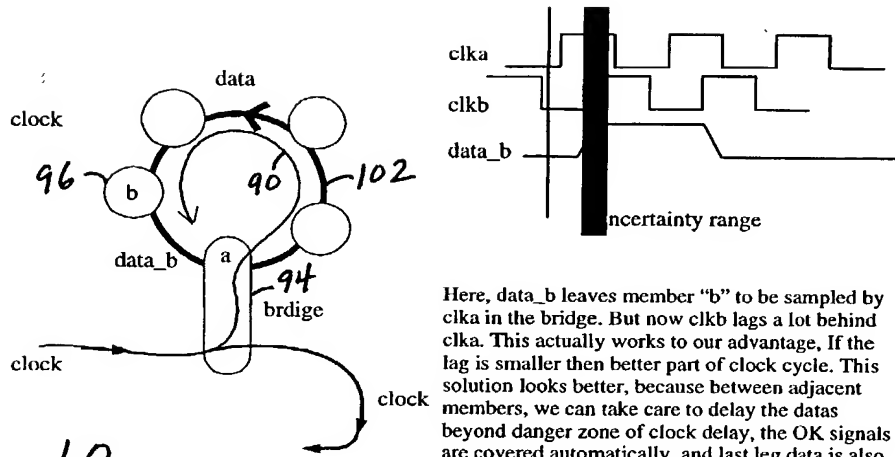


Fig. 10

Here, data\_b leaves member "b" to be sampled by clka in the bridge. But now clkb lags a lot behind clka. This actually works to our advantage. If the lag is smaller than better part of clock cycle. This solution looks better, because between adjacent members, we can take care to delay the data beyond danger zone of clock delay, the OK signals are covered automatically, and last leg data is also covered. The only signal not safe is the OK from bridge to "b" member. It will need a latch in "b".





Fig. 11

Fig. 12



if module "a" sends a message to module "b", ring works fine. However if most of the traffic is from "c" to "b", this is more expensive in terms of latency.

Another problem is "peak latency". Suppose that, "a" transmits mostly to "d" and "b" mostly to "c". In this case communication between "b" and "c" suffers degradation in case that peak traffic coincide.






Fig. 13




Land bridge gets its name from the fact that it is a luxury. It spans across connected modules. The idea is simple. When V2 sends message to D1 it gets to one side of the bridge. This side analyzes the destination address and by some magic (explained later) decides to short-cut the path. The message re-appears at the other end of the bridge and gets fast to D1. By same magic, message from V1 to D2 get bypassed also. message from V1 to D1 is treated directly.



Enumeration is started by "Anchor" which assigns address=1 to itself. results of enumeration are labels 1 to 7. land bridge gets two addresses, as if it were not one module. there is "near" end, that got enumeration label "3", and the "far" end marked 6.

Fig. 15






msg1 and msg2 arrive at the same time.  
the bridge end must make a decision  
which message to forward first.

It can be shown that unwise decision can  
lead to freezout, deadlock and option price  
dropping to 5\$.

Therefore MSG2 gets the priority.

Fig.  
17

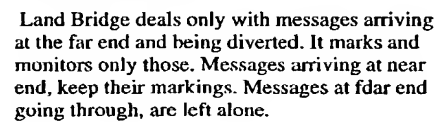



Fig. 19

Fig. 20




Fig. 21



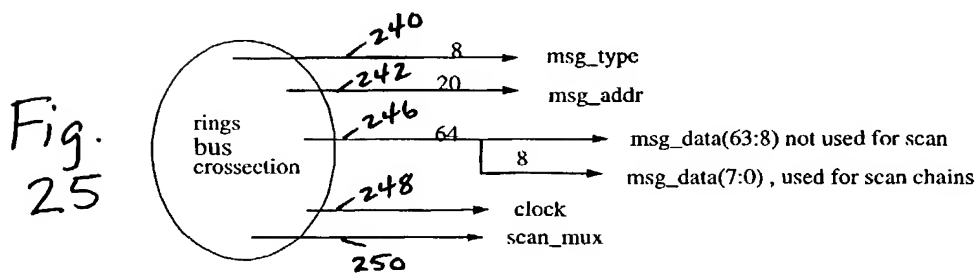




Fig. 22



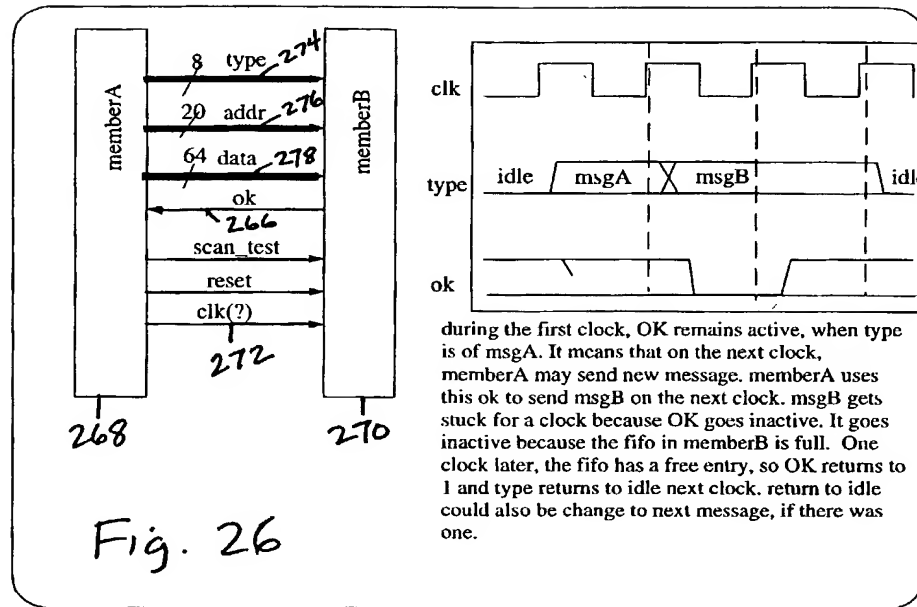


Fig. 27





Fig. 28



The incoming messages are examined first to see if it is supervisor or work/program. Work/program messages have address field. We check if it is our address. Since we know that our address is aligned to our power of 2, The address mask (named split mask) causes only certain number if upper bits to be compared. The lower part of the address is passed inside as internal address. The upper bits are compared against self-address register. This register gets its value during enumeration protocol. The lower part of this register is always masked. Hopefully synthesis will delete the unused bits implementation.




Fig. 29

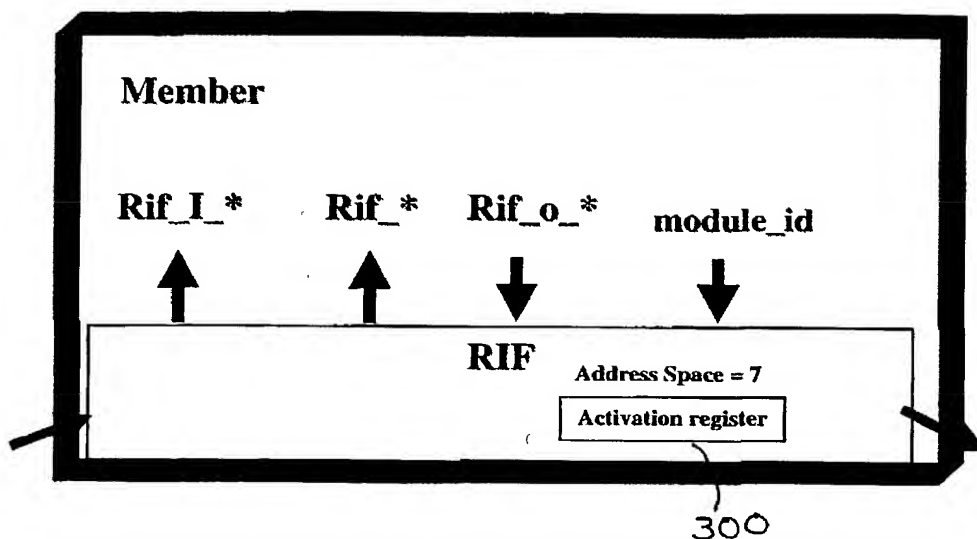


Fig. 30



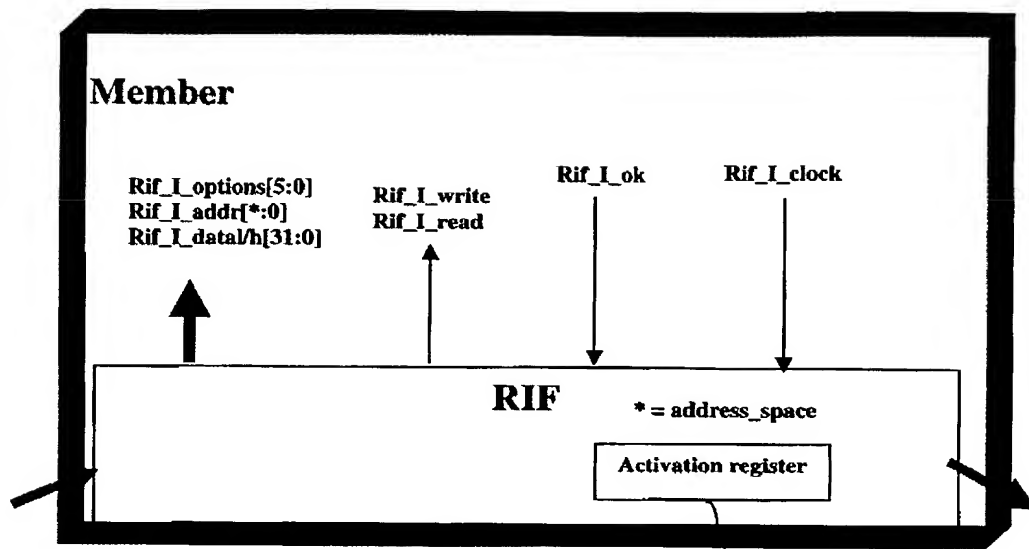


Fig. 31

300

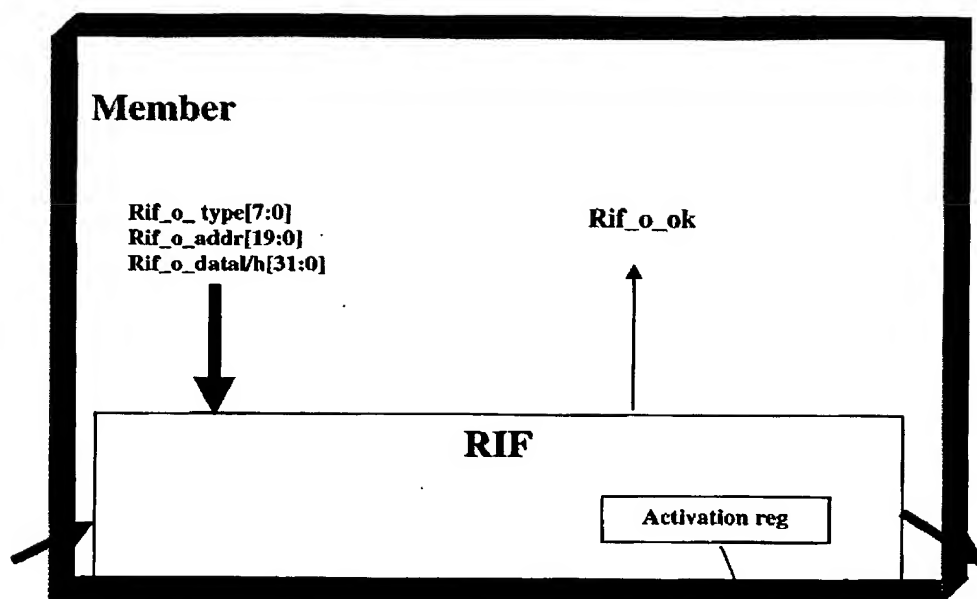


Fig. 32

300

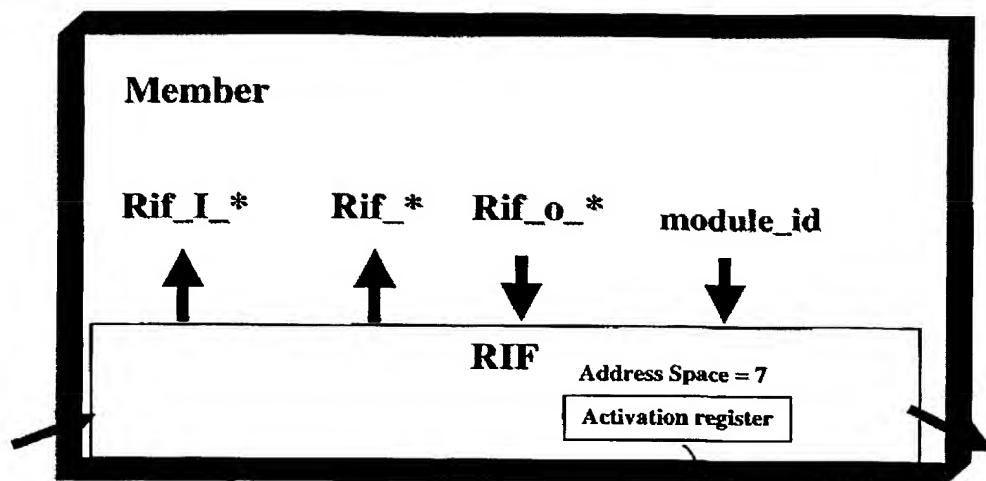



Fig. 33


300

Fig. 34




the second land bridge solves most traffic problems, but adds 4 clocks in the overall ring length. This is not a big problem because no message should travel the whole perimeter.

Fig. 35



The utopia interface is forced into mode that communicates in messages, not cells. We using the I/O and maybe some of the logic.

Fig. 36



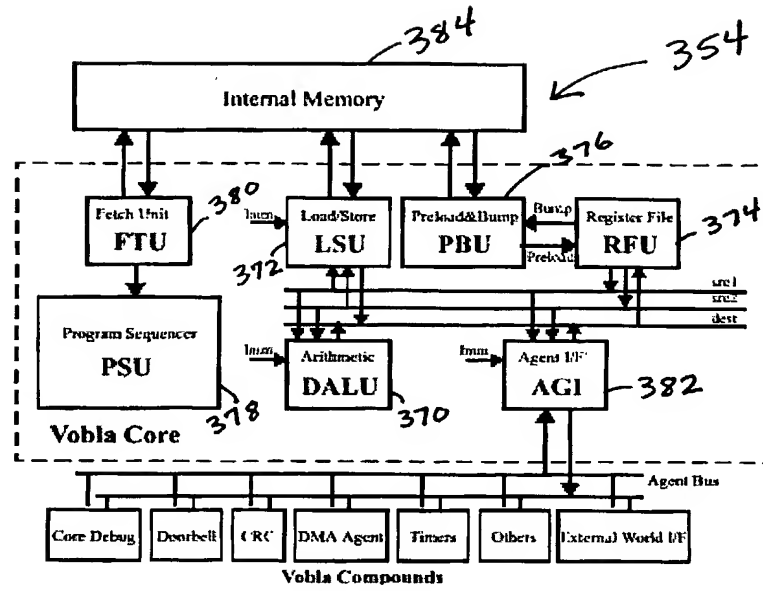



Fig.  
37

Fig.  
38



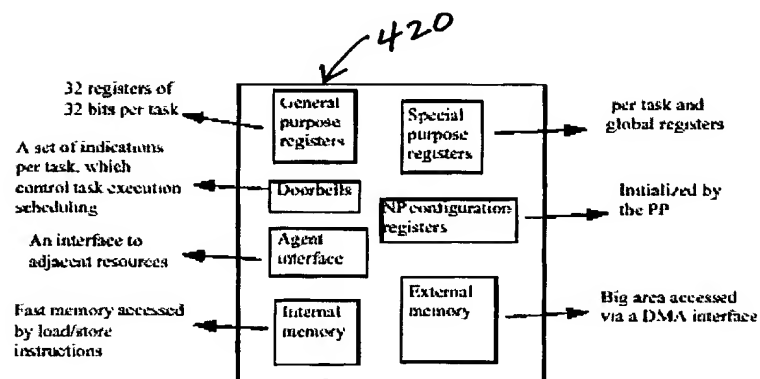
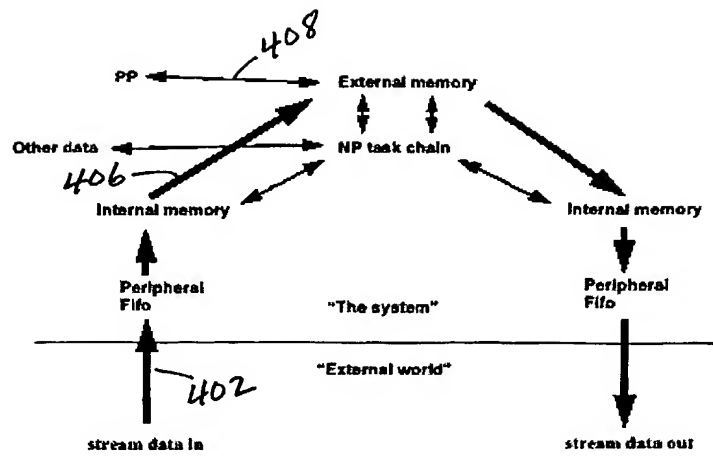


Fig. 41

R1 register:



s - sticky bit  
 eq - equal/zero  
 lt - less then/negative  
 gt - greater then/positive  
 c - carry  
 mb - reflection of the RAM multi-reader busy indication.

430


REFETCH SPR  
(spr index = 0)TASK SPR  
(spr index = 1)TRAP SPR  
(spr index = 2)MINDEX SPR  
(spr index = 3)

Fig. 42

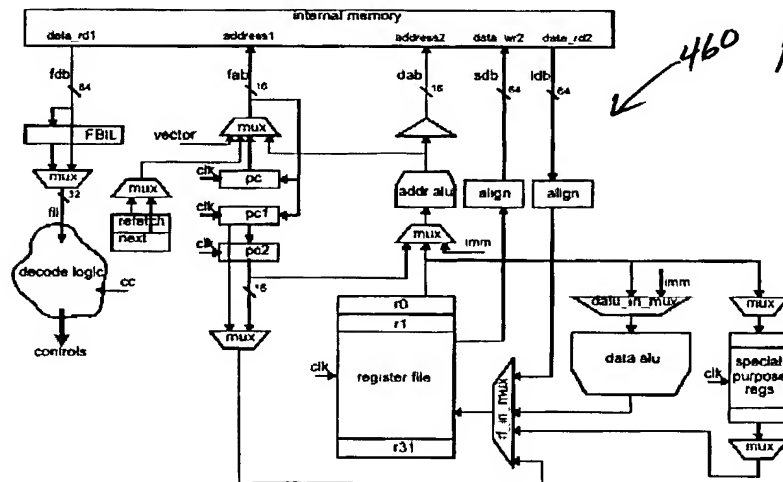
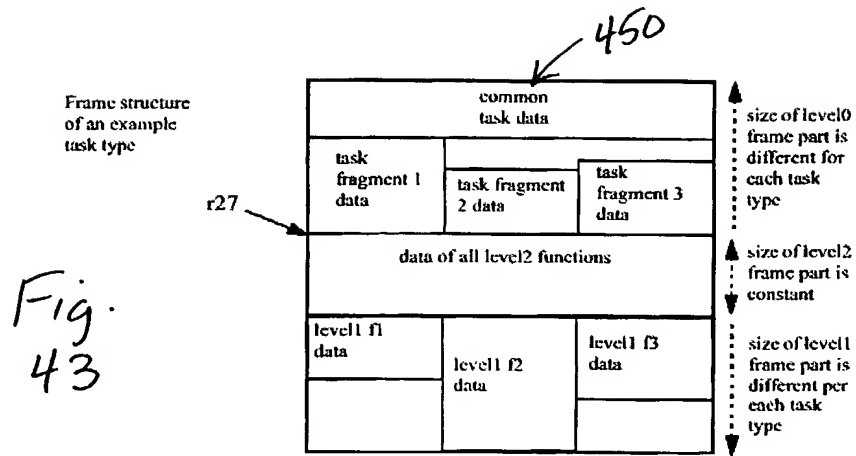
440

442

444

446





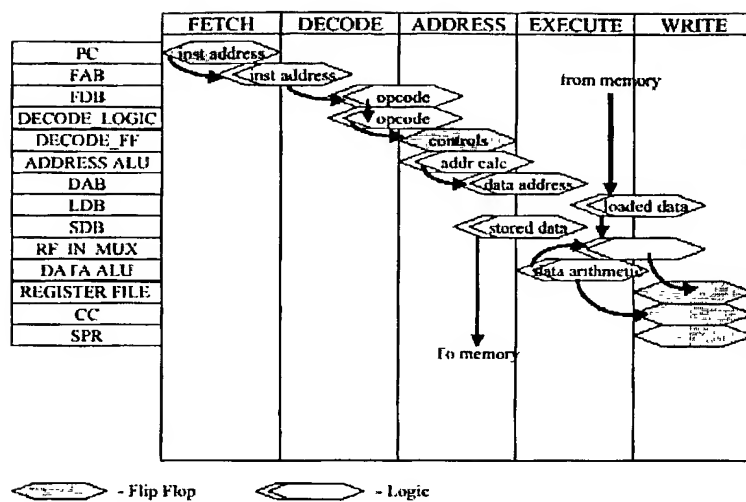
Fig.  
45

Fig. 46

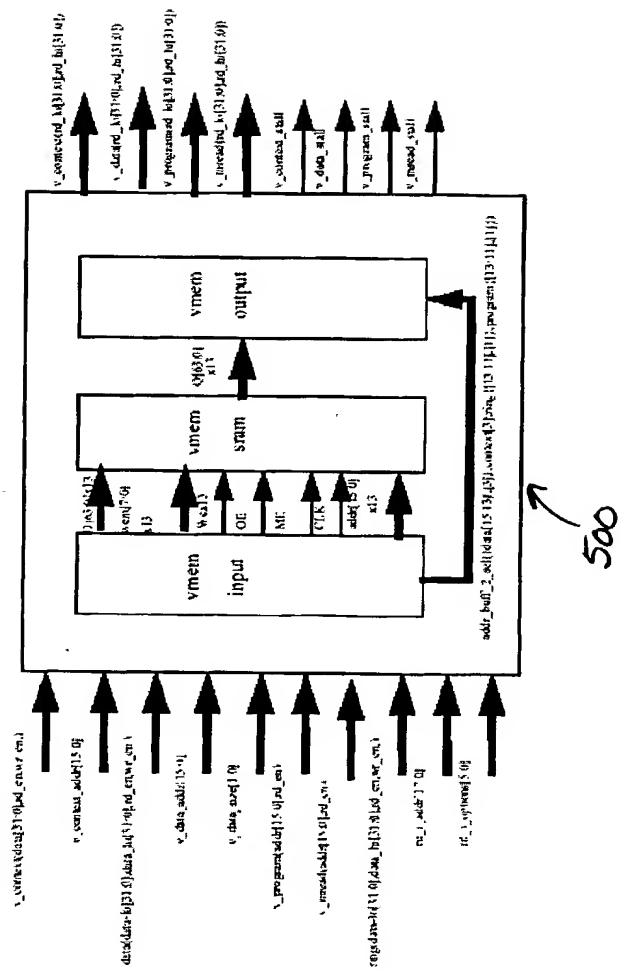
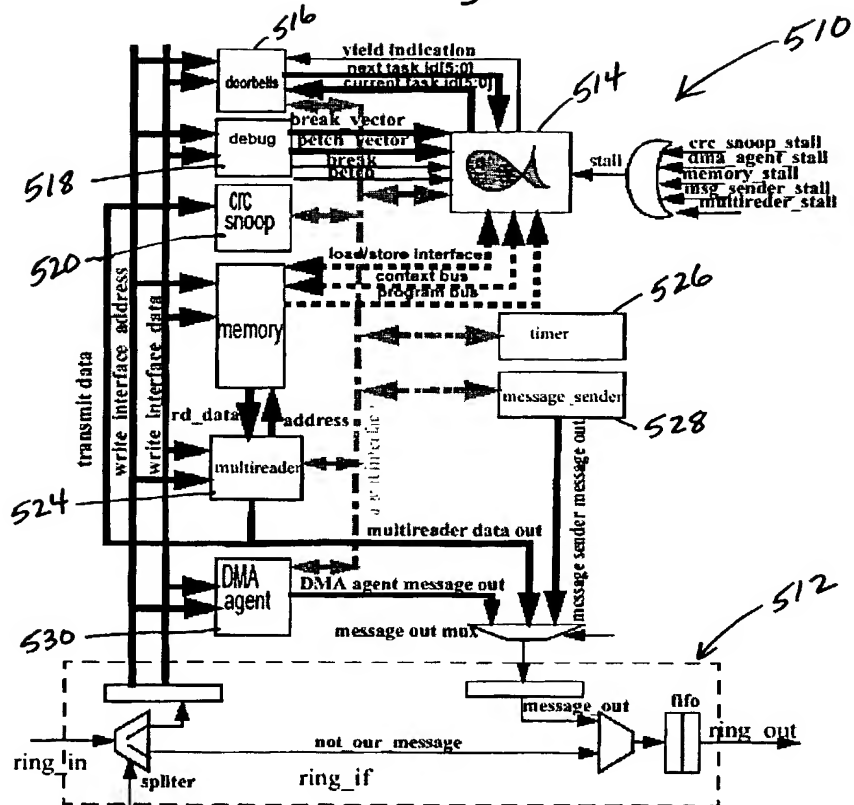


Fig. 47



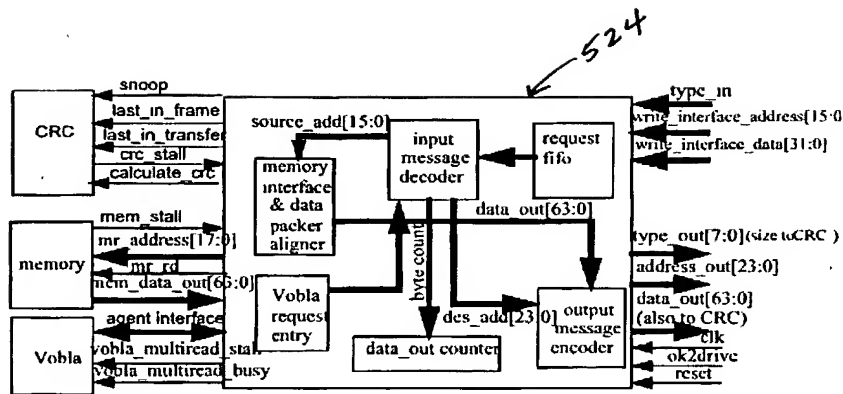


Fig. 48




Fig. 49


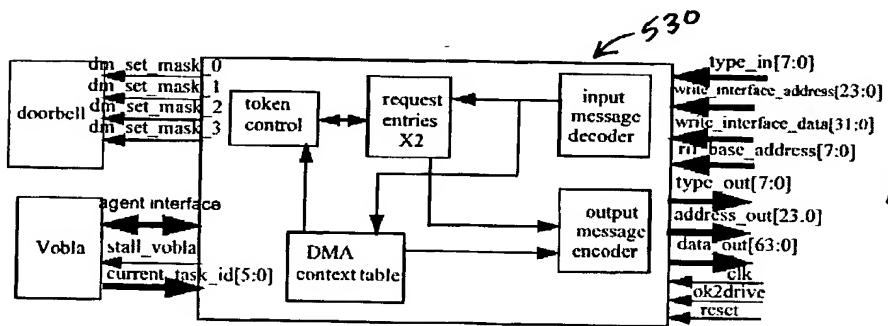
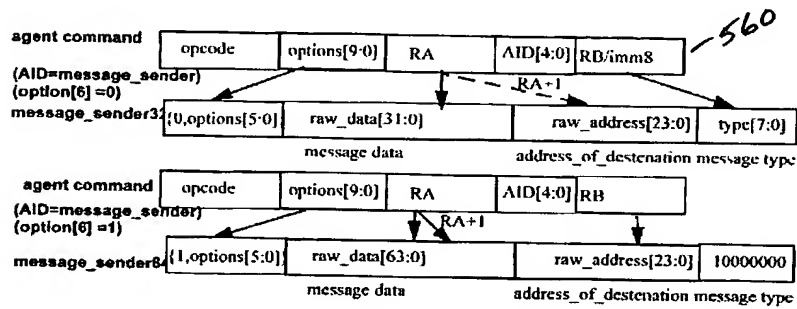
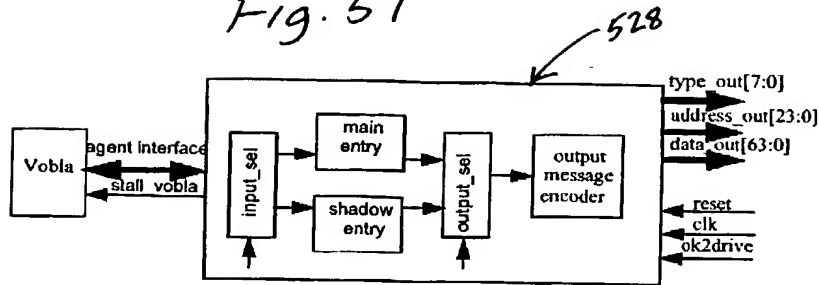


Fig. 50

Fig. 51





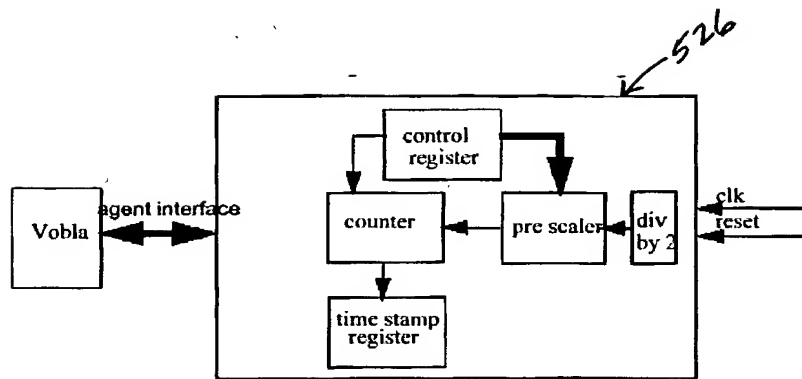


Fig.  
57

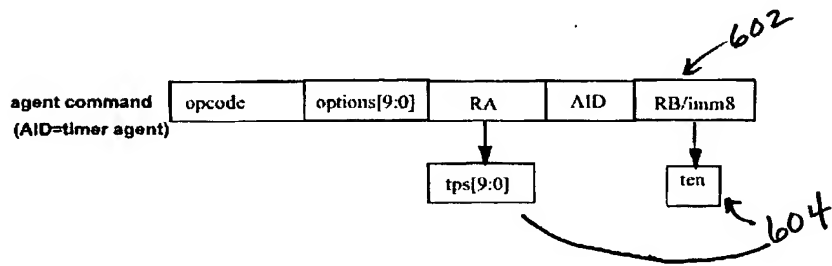






Fig.  
58












Fig. 64

Fig. 65





Fig. 66

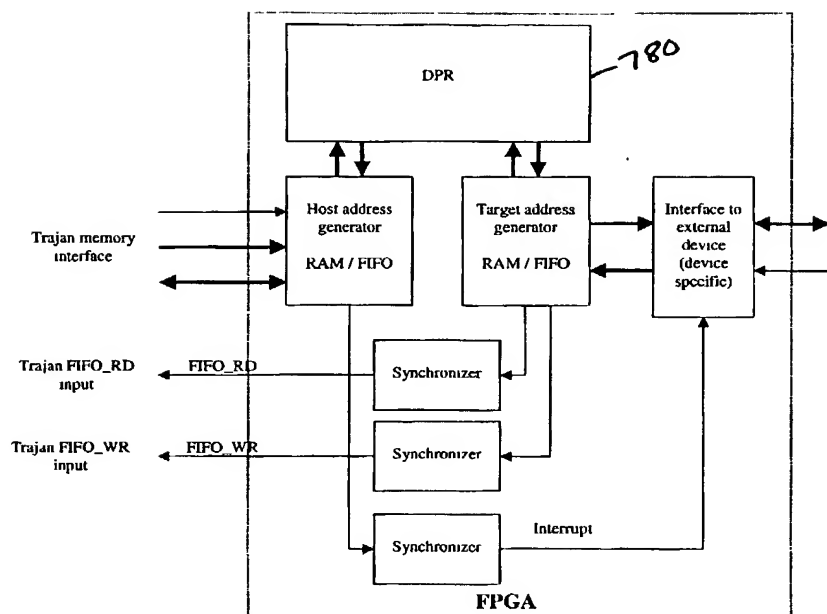


Fig. 67

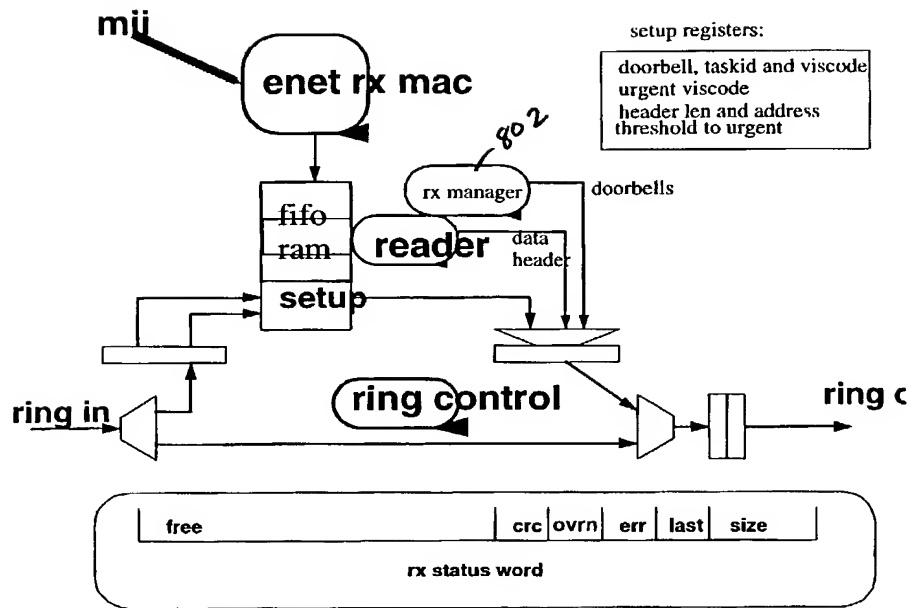



Fig. 68





### Data Plane

A block diagram of a network node architecture. At the top, a horizontal bar represents a network interface, with a handwritten label '830' and an arrow pointing to it. Below this is a large rectangular block representing the node. Inside this block, at the top, is a box labeled 'Protocol Processor'. Below the 'Protocol Processor' is a horizontal oval labeled 'Network Interface'. Below the 'Network Interface' are two boxes, 'Network Processor' on the left and 'Network Processor' on the right. To the left of the 'Network Processor' box is a vertical bar labeled 'Memory Interfaces'. To the right of the 'Network Processor' box is a vertical bar labeled 'peripherals'. A handwritten label '834' with an arrow points to the left 'Network Processor' box, and a handwritten label '836' with an arrow points to the right 'Network Processor' box. A large arrow points from the left into the 'Memory Interfaces' bar.

Fig  
70

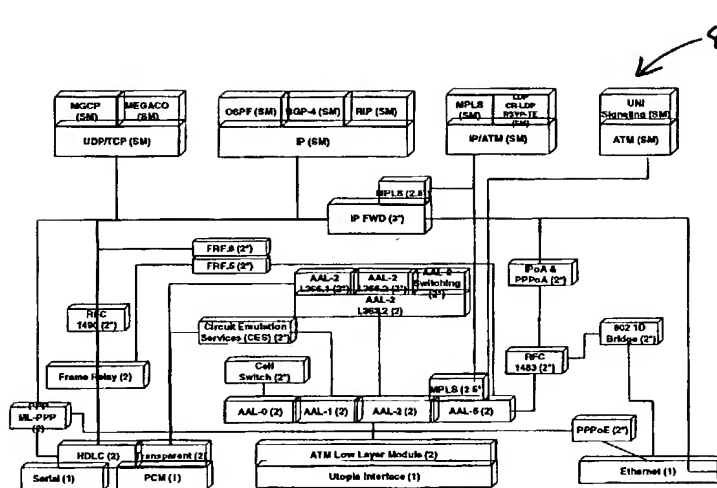


Fig.  
71




Fig. 72

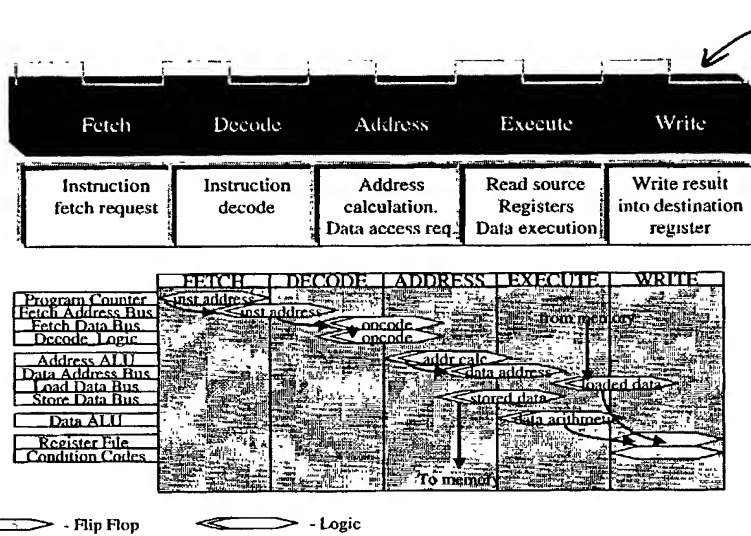
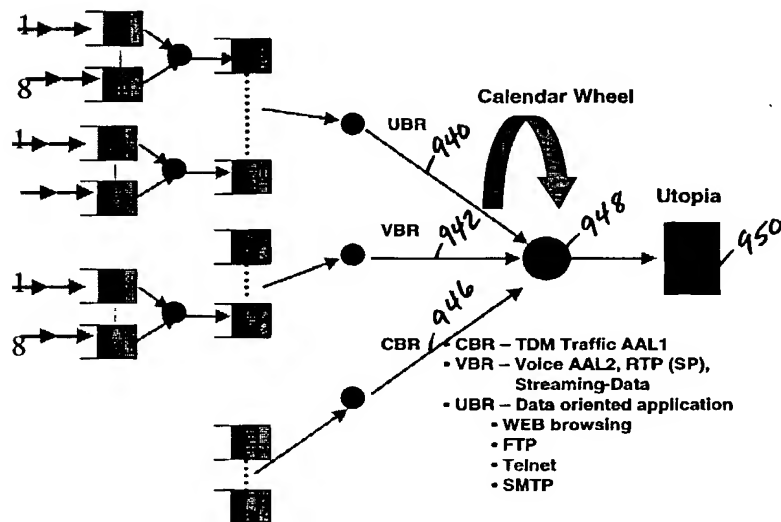


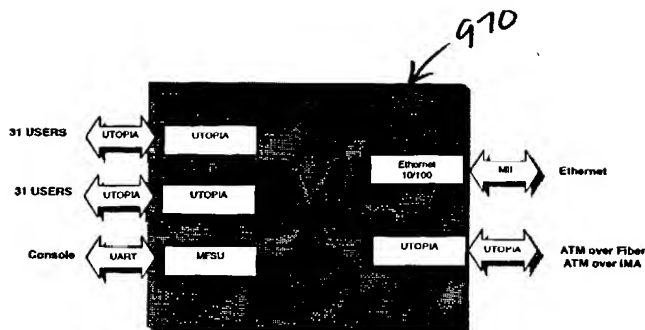


Fig. 73



Fig.  
76Fig.  
77

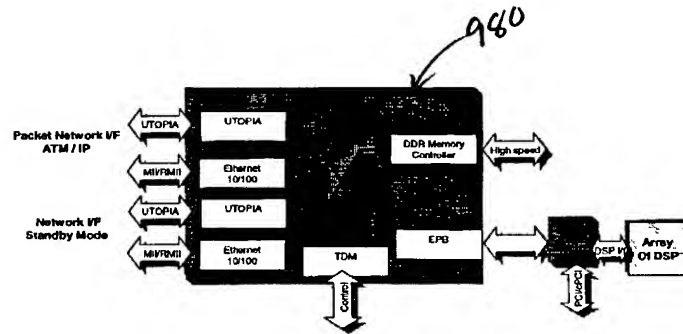


Fig. 78

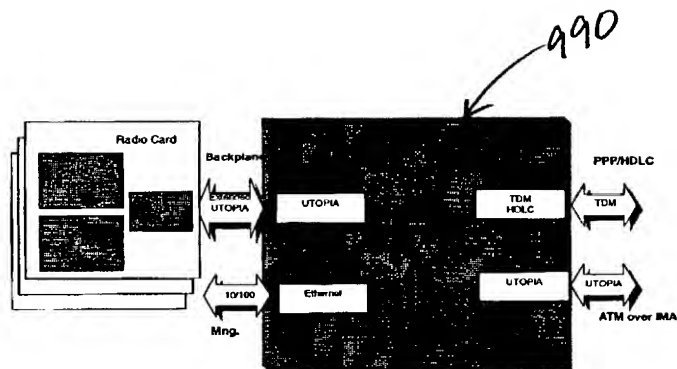


Fig. 79

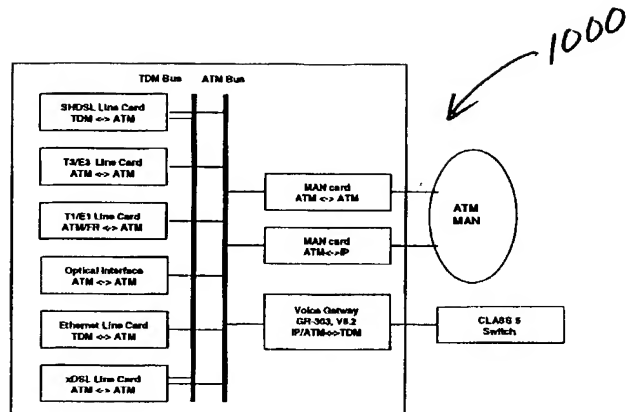


Fig. 80

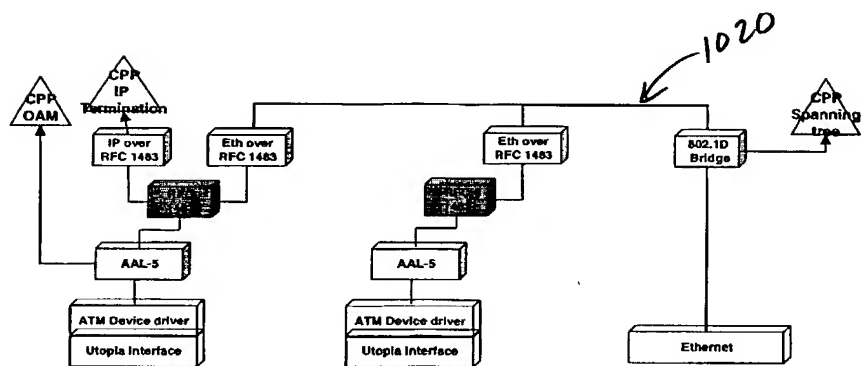


Fig. 81




Fig. 82

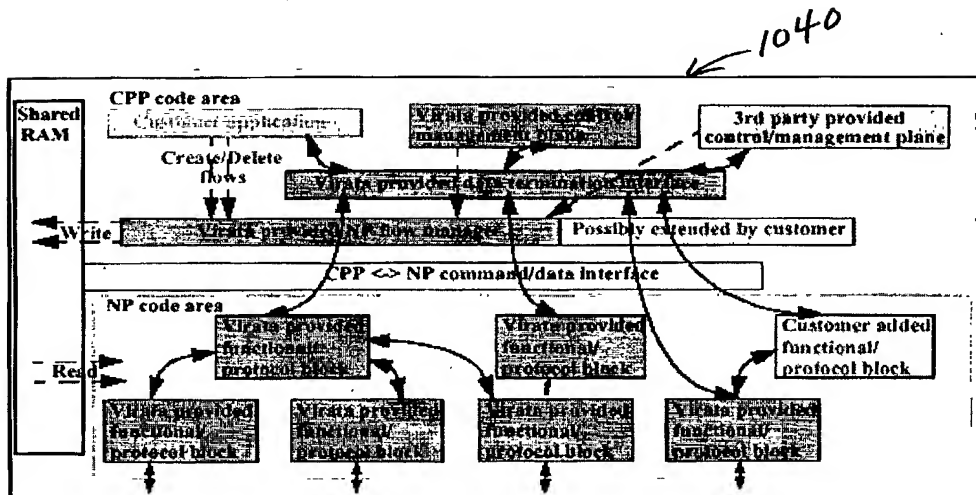


Fig. 83

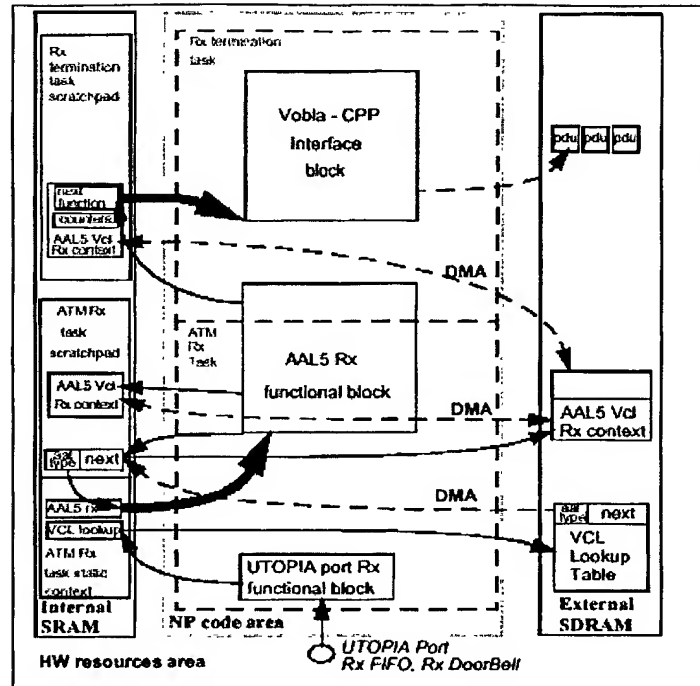
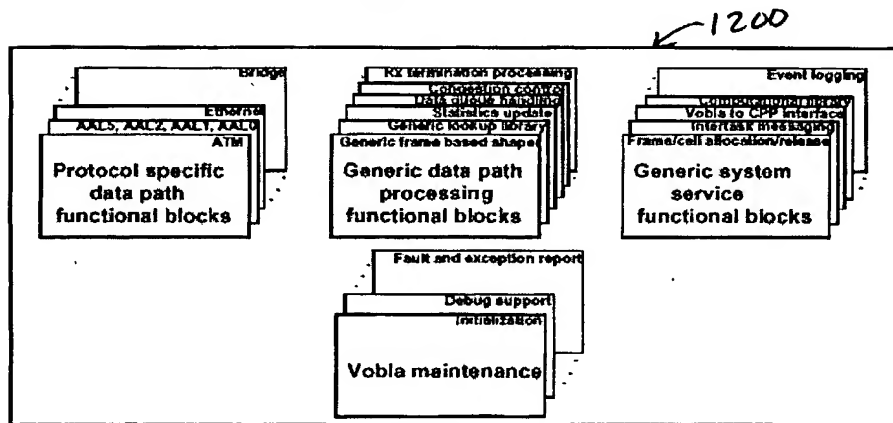



Fig. 84





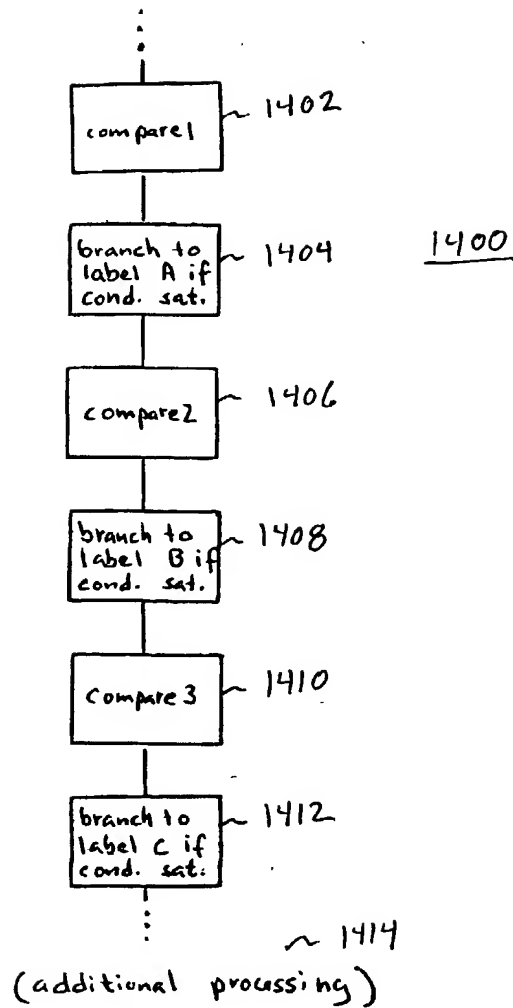


FIG. 87  
(PRIOR ART)

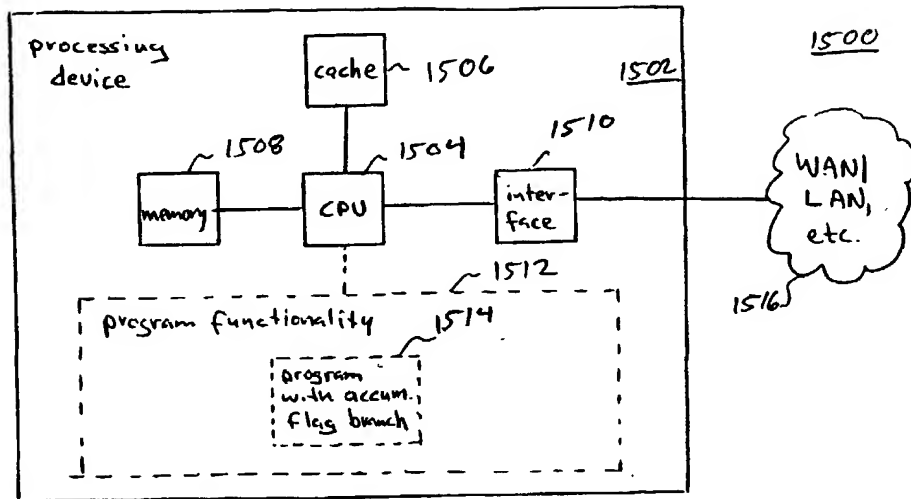


FIG. 88

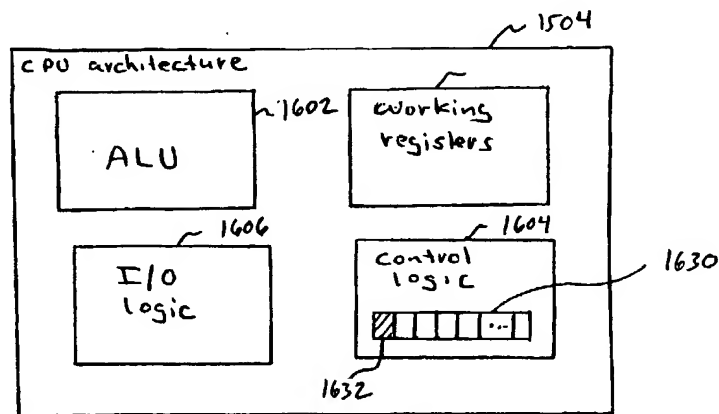


FIG. 89

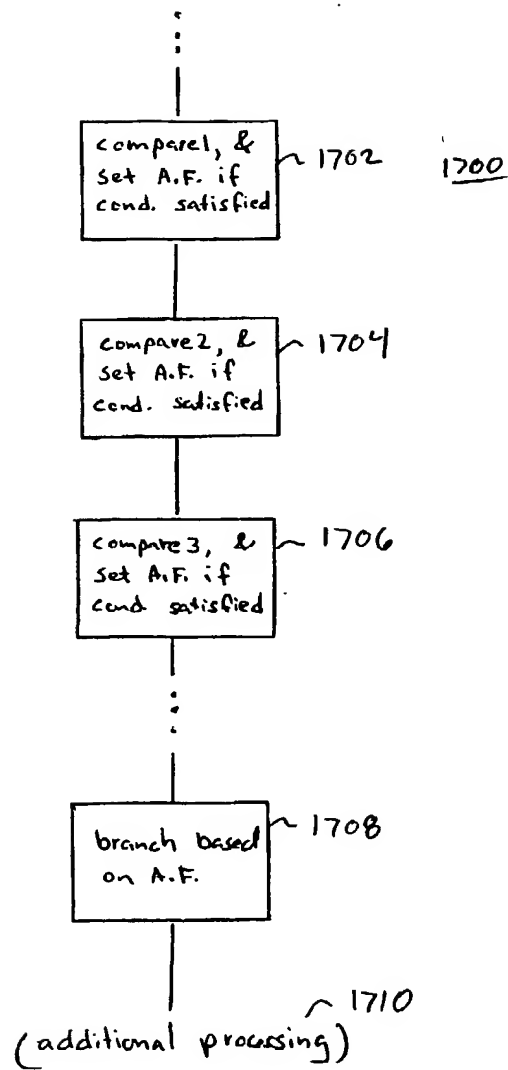


FIG. 90